PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 : H04H 1/00	A1	(11) International Publication Number: WO 93/11617 (43) International Publication Date: 10 June 1993 (10.06.93)

(21) International Application Number: PCT/US92/09785

(22) International Filing Date: 20 November 1992 (20.11.92)

(30) Priority data:
797,298 25 November 1991 (25.11.91) US

(71) Applicant: ACTV, INC. [US/US]; 1270 Avenue of the Americas, Suite 2401, Rockefeller Center, New York, NY 10020 (US).

(72) Inventor: FREEMAN, Michael, J. ; ACTV, Inc., 1270 Avenue of the Americas, Suite 2401, Rockefeller Center, New York, NY 10020 (US).

(74) Agents: DALEY, Christopher, J.; Dorsey & Whitney, 1330 Connecticut Ave. N.W., Suite 200, Washington, DC 20036 (US) et al.

(81) Designated States: AU, BG, BR, CA, CS, FI, HU, JP, KP, KR, LK, PL, RO, RU, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, SE).

Published
With international search report.
With amended claims and statement.

(54) Title: COMPRESSED DIGITAL DATA INTERACTIVE TELEVISION SYSTEM

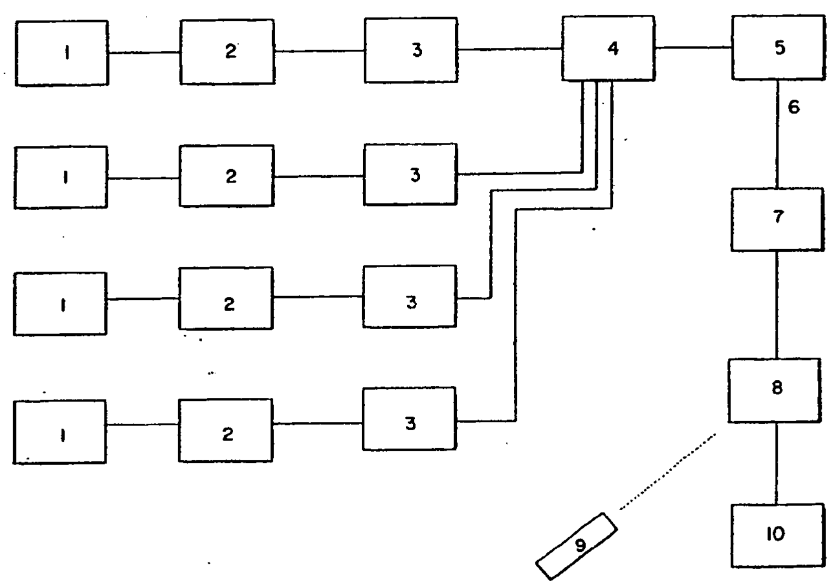

(57) Abstract

An interactive cable television system is disclosed which utilizes a standard television cable distribution network (6) for simultaneously providing a plurality of viewers with an interactive television program comprising a plurality of signals related in time and content. The video signals are transmitted in a digital format (2), more than one signal being multiplexed (4) into a data stream for transmission of multiple signals over a single channel. The digital video signals may be compressed (3) for transmitting more video signals per channel. A receiver (7), in conjunction with a signal selector (8), selects a particular NTSC channel for playback, then selects a particular video signal from the multiplexed signal, and uncompresses the video signal for playback to a television monitor (10). An alternative embodiment is disclosed wherein the various signals which comprise the interactive program are switched between at the head end rather than at the receiver. The multiple choice control unit (9) selects a desired signal by relaying the multiple choice selections of the user through a relay box back to a remotely located switching station (4). The switching station routes the correct video signal down the appropriate cable channel for the particular user.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	FR	France	MR	Mauritania
AU	Australia	GA	Gabon	MW	Malawi
BB	Barbados	GB	United Kingdom	NL	Netherlands
BE	Belgium	GN	Guinea	NO	Norway
BF	Burkina Faso	GR	Greece	NZ	New Zealand
BG	Bulgaria	HU	Hungary	PL	Poland
BJ	Benin	IE	Ireland	PT	Portugal
BR	Brazil	IT	Italy	RO	Romania
CA	Canada	JP	Japan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo			SE	Sweden
CH	Switzerland	KR	Republic of Korea	SK	Slovak Republic
CI	Côte d'Ivoire	KZ	Kazakhstan	SN	Senegal
CM	Cameroon	LI	Liechtenstein	SU	Soviet Union
CS	Czechoslovakia	LK	Sri Lanka	TD	Chad
CZ	Czech Republic	LU	Luxembourg	TG	Togo
DE	Germany	MC	Monaco	UA	Ukraine
DK	Denmark	MG	Madagascar	US	United States of America
ES	Spain	ML	Mali	VN	Viet Nam
FI	Finland	MN	Mongolia		

COMPRESSED DIGITAL DATA INTERACTIVE TELEVISION SYSTEM

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to interactive response systems, and more particularly to an interactive television system which provides interactive programming using compressed, digital data having more than one video signal per broadcast channel, or a multiplexed signal within a digital format.

2. Description of the Prior Art

Interactive systems are well known in the art. By synchronizing parallel tracks of an information storage media, and relating the content of the various tracks, it was found that interactive activity could be simulated. For example, commonly owned Freeman, U.S. Patent No. 3,947,972 discloses the use of a time synchronized multi-track audio tape to store educational conversations. One track is employed to relay educational interrogatories to a user, and the remainder of the tracks, selectable by a switching mechanism, are used to convey responsive messages.

These systems progressed to interactive television, wherein multiple broadcast or cable channels were switched responsive to user selections to provide interactive operation. Commonly owned Freeman, U.S. Patent No. 4,847,700 discloses an interactive television system wherein a common video signal is synched to a plurality of audio channels to provide content related to user selectable responses.

Commonly owned Freeman, U.S. Patent No. 4,264,925 discloses the use of a conventional cable television system to develop an interactive system. Standard television channels with time synchronized content are broadcast to a plurality of users. Each user switches between channels responsive to interrogatories to provide interactivity.

These systems have been tailored to include memory functions so that the system can be more interactive, individually responsive, and so that customized messages may be given to the various categories of users responsive to informational queries. Freeman, U.S. Patent No. 4,602,279 discloses the use of a memory to store demographic profiles of television viewers. This information is stored to be recalled later for providing target specific advertising, for example.

These prior art interactive television systems were generally concerned with providing one signal (i.e. one video signal) per channel, whether the channel is on cable television, broadcast television, of VCR. Because cable and broadcast television channel capacity is becoming more limited as more and more cable channels are being utilized for conventional programming, and these systems require multiple channels, it would be desirable to reduce the channel capacity required for such systems while still providing at least the same level of interactivity. These disadvantages of the prior art are overcome by the present invention which provides an interactive television system which employs multiple, time-synchronized, content-related video signals per broadcast channel.

SUMMARY OF THE INVENTION

The present invention is an interactive cable television system which utilizes digital video signals to provide customized viewing responsive to user selections. A standard cable or direct broadcast satellite television distribution network is utilized for transmitting the interactive and other programming to users. The present invention allows plurality of viewers to be simultaneously provided with a plurality of different program information message signals related in time and content to each other. The interactive program comprises a plurality of video signals related in time and content to one another.

The video signals are converted into digital format for transmission. In a digital format, it is possible to transmit more than one video signal per cable television channel. Further, it is possible to transmit video signals via conventional telephone lines. If desired, the various digital video signals may be compressed before transmission. Compression allows an even larger number of video signal to be transmitted over a channel of the transmission media. A multiplexer combines the various digital signals into a reduced number of transmission data streams for transmission. The various NTSC television channels may be allocated in a predetermined fashion to maximize the number of simultaneously transmittable signals. The multiplexer in conjunction with the cable television transmission system multiplexes the desired video signals onto the desired channels, and transmits these signals over the NTSC channels. The number of video signals which may be multiplexed onto a single transmission channel will vary depending on the video signals to be transmitted. The television

channels containing multiplexed video signals are transmitted over a standard cable television distribution network, or direct broadcast satellite transmission system. A receiver receives the various television channels, some or all containing multiplexed or non-multiplexed digital video signals, and in conjunction with a signal selector, selects a particular channel for playback, then selects a particular video signal from the multiplexed signal, and finally expands the video signal, if necessary, for playback to a television monitor.

A multiple choice controller operates to control the receiver and signal selector to select a particular video signal for playback. The multiple choice controller may be programmed to map the different cable television channels, and the multiple signals thereon, to a serial numerical channel representation to simplify use by the user. The signal selector includes the necessary expansion apparatus corresponding with the compression scheme in use.

In practice, a user selects a desired interactive program to be viewed by selecting a cable or direct broadcast satellite television channel having multiplexed video thereon. using the multiple choice controller, the user selectably switches between the related video signals on the selected channel responsive to information displays or interrogatory messages, the signal selector de-multiplexing, expanding and displaying the selected signal.

If more signals were needed for an interactive program than were mappable to a single channel, the signal selector in conjunction with receiver may be programmed to switch between the various video signals as well as the various broadcast channels to provide the necessary level of interactivity.

The various information segments in the various video signals relate in real-time and content so that an interactive conversation can occur as the video signal is played back and the user responds to the various interrogatories on the video signals. The use of multiple signals per channel may be used for many types of interactive programs, including those disclosed in the previously mentioned U.S. Patents, for example, field synchronized multiple camera angles from a sporting event, or an interactive game show.

In a two-way embodiment, the various signals which comprise the interactive program are switched between at the head end rather than at the receiver. This embodiment may be used for example in a cable television system, a

direct broadcast satellite system, a conventional telephone system modified to receive digital video signals, or any other appropriate transmission system capable of sending digital video signals. The multiple choice control unit, rather than selecting a desired signal from the a group of incoming signals, selects a desired signal by relaying the multiple choice selections of the user through a relay box back to a remotely located switching station, preferably the cable television source. The multiple choice selections may be relayed to the switching station in any conventional means, such as two-way cable television, telephone, or FM transmission. If the interactive programming is being transmitted over a telephone line, the multiple choice selections may be relayed back over the same telephone line. The switching station receives the multiple choice selection of the user and routes the correct signal down the appropriate cable channel, telephone line, or other transmission media for the particular user. In such an arrangement, only a single link is required between the subscriber or receiver and the head end so that the one channel link can be used to receive a plurality of different channel selections dependent on the interactive choice relayed from the receiver to the video switch at the head end.

If desired, the two-way link may be used for other purposes, such as to transmit user demographic data back to the programming source for commercial reasons, or to allow an interactive game show player to win prizes, for example.

The system of the present invention allows improved performance by the compression algorithms in use. When a channel change has been requested by the user, a slight imperceptible delay is programmed to allow the expansion algorithm an opportunity to adjust to the rapid change from one video signal to another.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 is a block diagram of the Interactive Television System of the present invention.

FIGURE 2 is a block diagram of the system of the present invention in a two-way transmission configuration.

DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention is an interactive television system in which a plurality of viewers are simultaneously provided with a plurality of different program information message signals related in time and content to each other.

Preferably at a remote location from the viewer, a plurality of video signals 1 are provided, all related in time and content to one another. Video signals 1 may be, for example, various field and audio synchronized camera angles of a sporting event, or a game show having a content and host acting responsively to user selections. Alternatively, video signals 1 may be any video signals suitable for interactive conversation, such as those described in U.S. Patent Nos. 4,847,700, 3,947,972, 4,602,279, 4,264,925, 4,264,924, for example, the contents of which are incorporated specifically herein by reference. However, it is readily foreseen that various types of time and content related video signals exist which are suitable for interactive operation.

In previous systems, these various signals would be transmitted to a receiver on separate channels, each requiring for example, a separate 6 Mhz NTSC channel, assuming the system is an NTSC system although any type of television transmission, such as PAL, etc. may be employed if desired. By the present invention, video signals 1 are directed to analog-to-digital ("A/D") convertors 2 which convert the various video signals into digital format for transmission. A/D convertors 2 may be of any conventional type for converting analog signals to digital format. It is readily foreseen that an A/D convertor may not be needed for each video signal 1, but rather fewer convertors, or even a single convertor might be capable of digitizing the various video signals 1. It is further foreseen that interactive video programs might be delivered to a cable or other distribution network in pre-digitized and/or precompressed format. In a digital format, it is possible to transmit the various video signals over fewer transmission channels than if the video were in analog format.

The digital conversion results in very large amounts of data. It may therefore be desirable to reduce the amount of data to be sent, allowing thereby more signals to be sent over a single transmission channel. For example, a single frame of NTSC video represents over 350 Kbytes of data. Therefore, two hours of standard video is about 80 Gbytes. Since there are 30 frames/sec in such video, the data transfer rate is 22 Mbytes/sec. Such large amounts of data are difficult to process using current computer technology. However, it is foreseen that rapid advances in computerization will eventually permit reception of video at data rates sufficiently high to allow reception of uncompressed or expanded video in household systems.

In order to reduce the data transfer requirements, the various digital video signals may be compressed before transmission. The video may be compressed by any conventional compression algorithm, the two most common being "processor intensive" and "memory intensive."

The processor intensive approach performs compression by eliminating non-changing aspects of a picture from the processing in the frame-to-frame transfer of information, and through other manipulations of picture information involving mathematical computations that determine the degree to which a given motion or other in a picture is perceptible to the human eye. This approach depends on high-speed processing power at the transmission point.

The memory approach involves division of a picture frame into hundreds of minuscule blocks of pixels, where each block is given a code representing its set of colors and variations in luminance. The code, which is a much smaller increment of information than all the information that would describe a given block of the picture, is transmitted to the receiver. There, it calls up the identically coded block from a library of blocks stored in the memory of the receiver.

Thus, the bit stream represents a much smaller portion of the picture information in this approach. This system is generally limited by the variety of picture blocks which may be stored in the receiver, which relates directly to memory size and microprocessor power.

Data Compressors 3 are provided to reduce the data for each video signal which must be transmitted. Data compressors 3 may be of any conventional type commonly known in the art for compressing video images, such as those previously described. It is foreseen that compression of the various video signals might be done with fewer data compressors 3 than one compressor per video signal. In a conventional analog NTSC system, by way of example, it is possible to transmit one video signal per 6 MHZ channel. By digitizing the video signal, it is possible to send more than one video signal per channel. Compressing the digitized signals, allows even more video signals to be transmitted over a single transmission channel. The number of signals which may be sent over a single channel is generally related to, for example, a) the type of video being sent; b) the video compression scheme in use; and c) the current state-of-the-art in computer and memory power; and d) the bandwidth of the transmission channel.

Compression techniques exploit the fact that in moving images there is very little change from frame-to-frame. Editing out the redundancies between frames and coding just the changes allows much higher compression rates. The type of video which normally contains a great deal of high-speed movement, such as occurs at live sporting events, will, therefore, have the lowest compression rates. Movies, on the other hand, which normally have a lower frame rate and less frame-to-frame change than a line sporting event will achieve higher compression rates. Currently, compression can be varied from 2:1 to 10:1 for satellites, and 2:1 to 5:1 for cable television systems, depending on the degree of motion. However, it is readily foreseen that compression techniques will improve in the future to provide larger compression rates. It is further foreseeable that computer speeds may increase to a level of performance which will allow uncompressed or expanded video to be transmitted at more than one signal per channel.

Once the various video signals 1 have been digitized and compressed, multiplexer 4 combines the various digital signals into a reduced number of transmission data streams for transmission. For example, if 68 NTSC channels are available, and each channel is capable of transmitting either 4 digitized, compressed slow moving video signals (e.g. movies) or 2 digitized, compressed, high-speed video signals (e.g. sports) then the various NTSC channels must be allocated in a predetermined fashion to maximize the number of simultaneously transmittable signals. Multiplexer 4 receives the incoming compressed, digitized video signals and in a predetermined conventional fashion, in conjunction with transmitter 5, multiplexes the desired video signal onto the desired channels, and transmits these signals over the NTSC channels. It is readily foreseen that certain NTSC channels will contain only one video or other signal, in analog or digital form.

As indicated earlier, the number of video signals which may be multiplexed onto a single transmission channel will vary. The transmission data stream is transmitted by transmitter 4 via transmission media 6 to a receiving station 7. The transmitter 4, media 6, and receiver 7 may be any conventional means for transmitting digital video signals including broadcast television, cable television, direct broadcast satellite, fiber optic, or any other transmission means. The transmission means may even be a telephone system capable of transmitting a digital video data stream. Thus, a multiplexed data stream having several channels

may be sent directly to a user over a single telephone line. It is readily foreseen that the aforementioned digital transmission devices may include means for transmitting analog signals as well.

In a preferred embodiment, the digital transmission signal is preferably transmitted via cable television. Receiver 7, receives the various NTSC channels, some or all containing multiplexed or non-multiplexed digital video signals. Ordinarily, more than one channel will be transmitted by transmitter 5 and received by receiver 7 as in an ordinary cable television system. However, each of the different channels may have several digitized video signals thereon. Therefore, receiver 7 preferably operates in conjunction with signal selector 8 to select a particular NTSC channel for playback, to select a particular video signal from the multiplexed signa and finally to uncompress or expand the compressed video signal, if necessary for playback to monitor 10.

Multiple choice controller 9 operates to control receiver 7 and signal selector 8 to select a particular video signal for playback. In practice, a user need not know that multiple signals per channel are in use. If, for example, 68 channels with 4 signals-per-channel were in use, controller 9, in conjunction with receiver 7 and signal selector 8 might be programmed to represent these channels to the user as channels 1272. Output 10 is for example a conventional television. Signal selector 8 preferably includes a conventional de-multiplexer for selecting a particular signal from the channel currently being received by receiver 7. Signal selector 8 further includes the necessary un-compression or expansion apparatus corresponding with the compression scheme in use by compressors 3.

In practice, a user would select a desired interactive program to be viewed by selecting a cable television station or direct broadcast satellite station having multiplexed video thereon. Using multiple choice controller 9, the user selectably switches between the related video signals on the selected channel channels responsive to information displays or interrogatory messages, signal selector de-multiplexing, uncompressing or expanding and displaying the selected signal.

For example, an interactive sporting event program might be transmitted on a 6 MHZ cable television signal using a compression-multiplexing scheme which allows two sports channels to be transmitted over a single NTSC channel. It might be desired to have four video signals for the particular interactive sporting event. A

first video signal might contain the standard broadcast signal of the game; the second signal might contain a close-up view of the game action; a third signal might contain a continuously updated replay of game highlights; the fourth signal might contain statistical information. These four video signals might for example be multiplexed as follows: signals one and two multiplexed onto cable channel 34; signals three and four multiplexed onto cable channel 35. These four signals might, however, be mapped by controller 9 to playback as channels 78, 79, 80, and 81 for the user. Each video signal of this interactive program might then includes a label which reads, for example, "Full-Screen Action -- Press 78: Close-up Action --Press 79: Replay -- Press 80: Statistics -- Press 81."

As shown, if more signals were needed for an interactive program than were mappable to a single channel, signal selector 8 in conjunction with receiver 7 may be programmed to switch between the various video signals 1 as well as the various broadcast channels to provide the necessary level of interactivity.

The multiplexed interactive program might be transmitted over a single telephone line, if desired. In this embodiment, multiple choice controller 9 would be programmed to switch between the various video signals on the single telephone line. If additional channels were desired, a two-way configuration might be used as described below.

The system of the present invention may be utilized in an educational embodiment. Information is stored on each video signal in a plurality of reproducible information segments, each of which comprises a complete message reproducible by the receiver directly in response to the selection of the video signal by signal selector 8 responsive to a user selection on multiple choice controller 9. Each of the information segments in the various video signals 1 contain interrogatory messages with associated multiple choice responses, responsive messages, informational messages, or combinations thereof. The messages contained in the various video signals 1 may include responsive messages, informational messages, interrogatory messages or combination thereof whose contents are related in real-time to particular interrogatory messages, and correspond to the multiple choice selectable responses to the particular interrogatory messages.

The various information segments in the various video signals relate in real-time and content so that an interactive conversation can occur as the video signal is played back and the child responds to the various interrogatories contained in the video signals. As a child answers a particular interrogatory with a multiple choice response, the information in the video signal associated with the particular selection is played back by the signal selector 7. In the selected video signal at the time at which the selection occurred, is an information segment whose content corresponds with the selected response to the previous interrogatory, whether or not the interrogatory was in the same video signal as the information segment being output. The various interrogatories, responsive messages, and informational messages may generally be contained in any or all of the various video signals provided that they are synchronized properly so as to retain a timed relationship, and correspond properly a decision tree logic.

The use of multiple signals per channel may be used for many types of interactive programs, preferably those disclosed in the previously mentioned U.S. Patents. It is readily foreseen that other interactive programs may be developed which are within the scope of the present invention.

As shown in FIG. 2, the system of the present invention may be operated in a two-way configuration. In this mode, the various video signals 1 are processed as previously described, being digitized by A/D convertor 2 and compressed by video compressors 3. The signals are then routed to a central switching station 4. In this embodiment, the switching between the various video signals is accomplished at the head end rather than at the receiver. Multiple choice control unit 9 relays the multiple choice selections of the user through a relay box 7 back to the remotely located switching station 4. The multiple choice selections may be relayed by relay box 7 to the switching station via any conventional means, such as two-way cable television, telephone, or FM transmission, for example. Switching station 4 receives the multiple choice selection of the user and routes the desired signal to transmitter 5 which conventionally transmits the desired video signal down the appropriate cable channel for the particular user. If desired, transmitter 5 may also transfer conventional programming on the cable television channels not being used for interactive programming. Alternatively, switching station 4 may include multiplexing equipment as previously described, and thus operate multiple

interactive or noninteractive programs over a single television channel. However, a very large processing capability would be needed to operate in such a configuration.

For example, if it were desired to implement the interactive football game program as previously described, a single NTSC cable channel might be allocated for the program. However, in this instance, the four video signals would be present at the transmitting end. In response to a signal from wireless controller 9, a signal is sent by relay box 7 to the cable TV switching station which routes the desired video signal to the requesting viewer. Such a system requires very fast switching equipment, but is readily foreseeable using digital imagery.

Alternatively, it might be desirable to transmit the interactive sporting event over a single telephone line. When the user enters a selection on controller 9, a signal is sent via the telephone line to the central switching station which routes the desired signal of the interactive program over the user's telephone line so that a single link handles both the interactive choice being made at the receiver and the transmission of that choice, out of a plurality of choices, from the head end where the actual switching takes place in response to the interactive selection made at the receiver.

The two-way link between the user and the switching station may be used for other purposes. For example, demographic data may be transferred from the user to the broadcast network for commercial purposes, bills may be paid, a game show winner may be sent a prize, or other commercial or non-commercial purpose may be achieved.

As previously described, compression systems generally perform less efficiently when frame-to-frame content includes many changes in pixel content; i.e. during fast motion, or scenery changes. The system of the present invention may be advantageously programmed to ease the processing burden on the uncompression program. When a key on the controller is depressed to select a desired signal, a slight imperceptible delay might be effectuated if desired. This delay would allow the uncompression or expansion algorithm a short period of time to adjust to the rapid change from one video signal to another which ordinarily causes a degradation in its efficiency. Utilizing this delay, it may be possible to increase the number of signals which may be transmitted per channel.

Although the present invention has been described in detail with respect to certain embodiments and examples, variations and modifications exist which are within the scope of the present invention as defined in the following claims.

What is claimed is:

1. In an improved interactive television system having:

a plurality of television reception systems (7,10), each of said television reception systems (7,10) comprising a television receiver (7), each of said television receivers (7) having a plurality of different television reception channels, each of said television reception channels having a different communication frequency; and,

a television programming transmission means (5) for substantially simultaneously providing a multi-information television program signal to said television reception systems (7,10), said multi-information television program signal comprising a plurality of simultaneously provided different information signals (1) each related in real time and content to each other;

wherein the improvement comprises:

said information signals (1) being in a digital format, and

each of said plurality of information signals (1) having a communication frequency corresponding with one of said plurality of different television reception channels, more than one of said information signals (1) being capable of having the same communication frequency and being combined into a multiplexed program signal with at least one other of said more than one of said information signals (1), said more than one information signals (1) in the multiplexed program signal being substantially simultaneously transmittable;

any of said different television reception channels being capable of carrying a single information signal or a multiplexed program signal;

each of said television receivers (7) being capable of independently selectably receiving any of said plurality of information signals (1), each of said television receivers (7) comprising (i) a multichannel selection means (8) for selecting the television reception channel frequency to be received, and (ii) signal selection means (8) for selecting a particular information signal from the multiplexed program signal, each of said information signals (1) having a numerical channel representation in said multichannel selection means (8);

at least one of said information signals (1) further comprising video information displayable on said television receiver (7) corresponding to informational labels to be dynamically assigned to said numerical channel

representations for a particular multi-information television program, said television displayable informational labels being dynamically variable dependent on the content of said multi-information television program,

whereby flexible multi-information television programming may be provided with a reduced number of television channels.

2. An improved interactive cable television system according to claim 1 wherein said information signals (1) comprise a plurality of successive information segments, and at least a portion of said plurality of said information segments on said plurality of information signals (1) being content related in a decision tree relationship between successive individual segments and between information signals (1) whereby a stored accumulation program format may be received as said selectable multi-information television programming.

3. An improved interactive television according to claim 1 wherein said television programming transmission means (5) further comprising means for substantially simultaneously providing at least one regular broadcast television signal, having a communication frequency, to said television reception systems (7,10) on said frequency, said television receivers (7) being capable of receiving said at least one regular broadcast signal on a channel corresponding with said frequency.

4. An improved interactive television system according to claim 3 wherein said multiplexed program signal further comprises at least one regular broadcast television signal.

5. An improved interactive television system according to claim 1 wherein said television programming transmission means (5) provides said multi-information television program signal to said television reception system by means of a one way television signal distribution network (6).

6. An interactive television system having a plurality of television reception systems (7,10), each of said television reception systems (7,10) having a television receiver (7), each of said television receivers (7) having at least one television reception channel, each of said at least one television reception channels having a different communication frequency,

a television distribution network (6) for distributing television programming to said plurality of television reception systems, and

a television transmission means (5) operatively connected to said distribution network (6) for providing said television programming to said plurality of television reception systems (7,10), said television transmission means (5) for providing a multi-information television program signal as said television programming to said television reception systems, said interactive television system comprising:

multiplexing means (4), coupled to said transmission means (5), for providing a multiplexed multi-information television program signal;

digitizing means (2), coupled to said multiplexing means (4), for digitizing a plurality of multi-information segments, each multi-information segment comprising a plurality of simultaneously provided different information signals (1) related in real time and content to each other, multiplexed (4) at a communication frequency so as to become said multiplexed multi-information television program signal which may be substantially simultaneously reproducible;

each of said television receivers (7) being capable of independently selectably receiving information on any one of said at least one television reception channel dependent on the television reception channel selected;

said television reception systems further comprising a multi-information selection means (8,9) coupled between said distribution network (6) and said television receiver (7), said multi-information selection means (8,9) comprising signal selection means (8) coupled to said subscriber distribution network (6) for demultiplexing said multiplexed multi-information television program signal and selectively providing an output signal comprising only one of said related different information signals (1) for regenerating said output signal on said television receiver (7,10) in said one television reception system;

whereby viewers in said television system may independently selectably view any of said simultaneously provided different information signals, said multi-information television programming being received on a multiplexed reception channel.

7. An interactive television system according to claim 6 wherein said television programming further comprises a plurality of different regular television program signals along with said multiplexed multi-information television program signal, each of said different regular television program signals having a

communication frequency, more than one of said regular television program signals being capable of having the same communication frequency, each of said different regular television program signals being capable of being multiplexed with said at least a portion of said information signals (1), each of said different regular television program signals being directly selectably receivable by said signal selection means (8) and displayable on a corresponding television reception channel.

8. An interactive television system according to claim 6 wherein said distribution network (6) being selected from the group consisting of cable television, telephone, broadcast television, and direct broadcast satellite.

9. An interactive television system according to claim 6 wherein said selection means (8,9) comprises keyboard means (9) comprising a plurality of keys for selectably enabling converting of any one of said associated frequencies of said simultaneously provided different information signals (1) into said television reception channel associated frequencies dependent on the key selected.

10. An interactive television system according to claim 9 wherein said multi-information segments comprise video information displayable on said television receiver (7) corresponding to information labels to be dynamically assigned to said keys for a particular multi-information segment, said television displayable information labels being dynamically variable dependent on the content of said particular multi-informational segment.

11. An interactive television system according to claims 6 or 10 wherein at least a portion of said plurality of said information segments on said plurality of information signals (1) being content related in a decision tree relationship between successive individual segments and between information signals (1) whereby a stored accumulation program format may be received as said selectable multi-information television programming.

12. An interactive television system according to claim 6 wherein at least at least two of said simultaneously provided different information signals (1) comprise different field synchronized camera angles of the same event.

13. In an improved interactive television system having a plurality of television reception systems, each of said television reception systems comprising a television receiver (7), each of said television receivers (7) having a plurality of

different television reception channels, each of said television reception channels having a different communication frequency;

a television distribution network (6) for distributing television programming to said plurality of television reception systems; and

a television transmission means (5) operatively connected to said distribution network (6) for providing television programming to said plurality of television reception systems (7,10), said television programming transmission means (5) having means for substantially simultaneously providing a multi-information television program signal as said television programming to said television reception systems;

said multi-information television program signal comprising a plurality of simultaneously provided different information signals (1) related in real time and content to each other;

wherein the improvement comprises:

digitizing means (2), for digitizing said plurality of simultaneously provided different information signals (1);

switching means (4), coupled between said digitizing means and said transmission means (5), for selecting one of said plurality of simultaneously provided different information signals (1) responsive to a control signal corresponding to a particular television receiver for transmitting said selected signal to said particular television receiver, said control signal being received over a transmission media;

each of said television receivers (7) being capable of independently selectably receiving any of said plurality of information signals (1) dependent on the television reception channel selected,

each of said television receivers (7) comprising a selection means (8) for selecting the television reception channel associated frequency to be received, and

signal selection means (7,8) for generating and transmitting said control signal to said distribution network (6) for signalling said switching means (4) to transmit a particular information signal from said plurality of information signals (1) to said television receiver (7),

whereby flexible multi-information television programming may be provided with a reduced number of television channels.

14. An improved interactive television system according to claim 13 wherein said information signals (1) comprise a plurality of successive information segments, and at least a portion of said plurality of said information segments on said plurality of information signals (1) being content related in a decision tree relationship between successive individual segments and between information signals (1) whereby a stored accumulation program format may be received as said selectable multi-information television programming.

15. An improved interactive television according to claim 13 wherein said transmitted television programming further comprises at least one regular broadcast television signal having a communication frequency, said television programming transmission means (5) further comprising means for substantially simultaneously providing said at least one regular broadcast television signal to said distribution network (6) on said communication frequency, said television receivers (7) being capable of receiving said at least one regular broadcast signal on a channel corresponding with said frequency.

16. An improved interactive television system according to claim 13 wherein said distribution network (6) being selected from the group consisting of two-way cable television, two-way telephone, and two-way direct broadcast satellite.

17. An improved interactive television system according to claim 13 wherein said control signal transmission media being selected from the group consisting of telephone, cable television, FM transmission, and fiber-optic.

18. An improved interactive television system according to claims 13 or 14 wherein each of said information signals (1) having a numerical channel representation in said multichannel selection means;

at least one of said information signals (1) further comprising video information displayable on said multichannel television receiver (7) corresponding to informational labels to be dynamically assigned to said numerical channel representations for a particular multi-information television program, said television displayable informational labels being dynamically variable dependent on the content of a said multi-information television program.

19. An improved interactive television system according to claims 2, 11 or 14 wherein said displayable informational labels dynamically vary according to the successive decision tree selections to be made.

20. An improved interactive television system according to claims 1, 6 or 13 further comprising means (3), coupled to said transmission means, for compressing said digital information signals (1).

AMENDED CLAIMS

[received by the International Bureau on 23 April 1993(23.04.93); original claim 19 amended; other claims unchanged (1 page)]

14. An improved interactive television system according to claim 13 wherein said information signals (1) comprise a plurality of successive information segments, and at least a portion of said plurality of said information segments on said plurality of information signals (1) being content related in a decision tree relationship between successive individual segments and between information signals (1) whereby a stored accumulation program format may be received as said selectable multi-information television programming.

15. An improved interactive television according to claim 13 wherein said transmitted television programming further comprises at least one regular broadcast television signal having a communication frequency, said television programming transmission means (5) further comprising means for substantially simultaneously providing said at least one regular broadcast television signal to said distribution network (6) on said communication frequency, said television receivers (7) being capable of receiving said at least one regular broadcast signal on a channel corresponding with said frequency.

16. An improved interactive television system according to claim 13 wherein said distribution network (6) being selected from the group consisting of two-way cable television, two-way telephone, and two-way direct broadcast satellite.

17. An improved interactive television system according to claim 13 wherein said control signal transmission media being selected from the group consisting of telephone, cable television, FM transmission, and fiber-optic.

18. An improved interactive television system according to claims 13 or 14 wherein each of said information signals (1) having a numerical channel representation in said multichannel selection means;

at least one of said information signals (1) further comprising video information displayable on said multichannel television receiver (7) corresponding to informational labels to be dynamically assigned to said numerical channel representations for a particular multi-information television program, said television displayable informational labels being dynamically variable dependent on the content of a said multi-information television program.

19. An improved interactive television system according to claim 2 or 14 wherein said displayable informational labels dynamically vary according to the successive decision tree selections to be made.

STATEMENT UNDER ARTICLE 19

Claim 19 has been amended in response to "Box I" of the International Search Report. "Box I" stated that claim 19 was not drafted in accordance with the second and third sentence of Rule 6.4(a). As originally filed, Claim 19, a multiple dependent claim, depended from claim 11 another multiple dependent claim. Claim 19 has been amended so that it no longer depends from claim 11. Applicant apologizes for any inconvenience this error may have caused.

Applicant requests acceptance of claim 19 and that the International Search Report be "established in respect of claim 19."

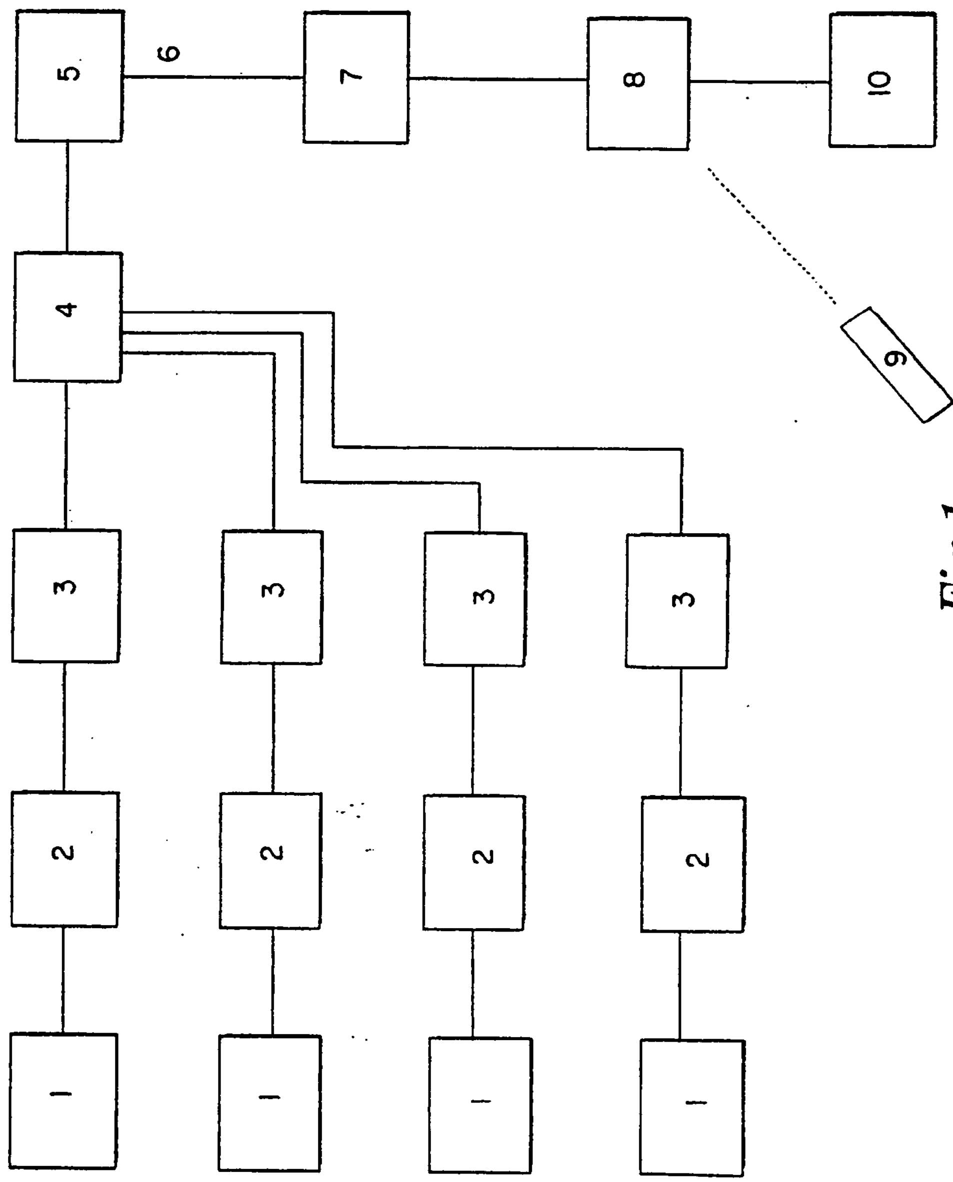

Fig.1

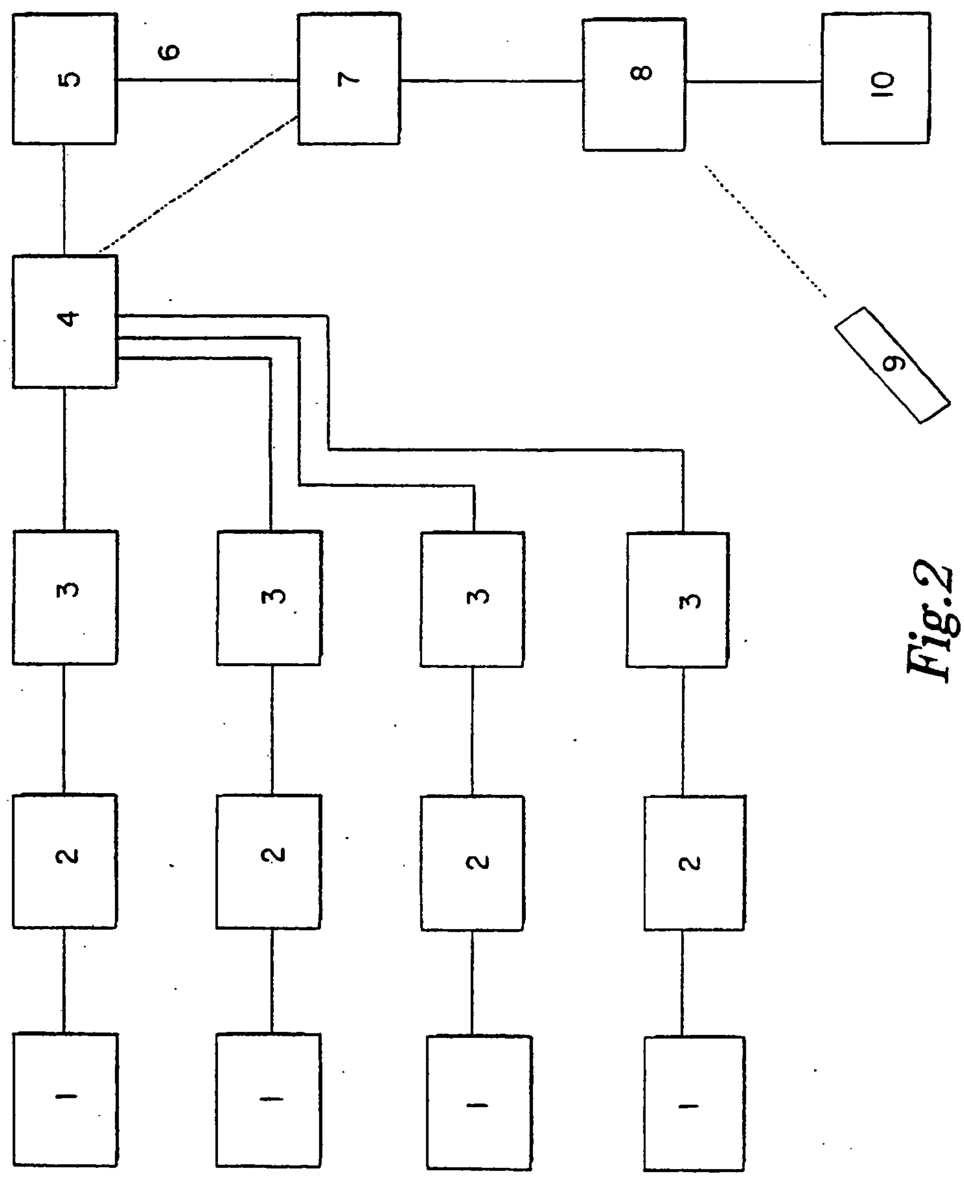

Fig.2

INTERNATIONAL SEARCH REPORT

PCT/US92/09785

A. CLASSIFICATION OF SUBJECT MATTER

IPC(5) :H04H 1/00
US CL :358/86; 455/3.1
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 455/4.1,4.2,5.1,6.1,6.2,6.3

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 4,264,925 (FREEMAN ET AL.) 28 April 1981, See columns 2-3, fig. 1.	1-18,20
Y	US, A, 4,975,771 (KASSATLY) 04 December 1990, See columns 2-8, figs. 1,3,6,8.	1-18,20
A,P	US, A, 5,133,079 (BALLANTYNE ET AL.) 21 July 1992, See columns, 1-4, fig. 1B.	1-18,20
A	US, A, 4,264,924 (FREEMAN) 28 April 1981, See columns 2-6, fig. 1.	1-18,20
A,P	US, A, 5,132,992 (YURT ET AL.) 21 July 1992, See columns 1-3, fig. 2A.	1-18,20

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
07 JANUARY 1993	25 FEB 1993
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. NOT APPLICABLE	Authorized officer Ngoc-Ho Nguyen for NGUYEN VO NGUYEN NGOC-HO INTERNATIONAL DIVISION Telephone No. (703) 305-4786

Form PCT/ISA/210 (second sheet)(July 1992)★

Multimedia Systems (1995) 3: 137–149

Multimedia Systems

P.D. 1995
p. 137–149 = 13

XP 000576898

Design and analysis of a look-ahead scheduling scheme to support pause-resume for video-on-demand applications

Philip S. Yu, Joel L. Wolf, Hadas Shachnai

IBM Research Division, Thomas J. Watson Research Center, Yorktown Heights, NY 10598, USA

Abstract. In a video-on-demand (VOD) system, subscribers can choose both the movie they wish to view and the time they wish to view it. In such an environment there are invariably "hot" videos which are requested by many viewers. The requirement that each viewer be able to independently pause the video at any instant and later resume the viewing with little delay can cause difficulties in batching viewers for each showing. Under batching, a single video stream is shared by multiple concurrent viewers and a resume request has to wait for additional stream capacity to become available before actual resumption can occur. The conventional approach to the support of on-demand pause-resume provides one video access stream to disks for each video request. This can greatly increase the disk arm requirements of a VOD system. In this paper, we propose a more efficient mechanism to support the pause-resume feature using *look-ahead scheduling* with *look-aside buffering*. The idea is to use buffering to increase the number of concurrent viewers supportable. The concept of look-ahead scheduling is not to back up each viewer with a real stream capacity so he can pause and resume at any time, but rather with a (look-ahead) stream that is currently being used for another showing which is close to completion. Before the look-ahead stream becomes available, the pause and resume features have to be supported by the original stream through (look-aside) buffering of the missed content. It is shown via simulations that the proposed scheme can provide a substantially greater throughput than the approach without batching. Furthermore, for a given amount of buffer, the improvement in throughput grows more than linearly with the stream capacity of the server. In other words, the look-ahead stream scheduling scheme operates with good economy of scale because it is easier to form look-ahead streams for video servers with larger stream capacity.

Key words: Video-on-demand – Pause-resume – Performance analysis

Correspondence to: P.S. Yu

1 Introduction

Multimedia servers are quite different from those of conventional computer file systems. For one thing, multimedia information, including motion video and audio, requires a guaranteed transfer rate. In NTSC video, for example, retrieval from or update to a storage system must have a sustained and almost constant rate of 30 frames/s. Besides the strict timing requirement, multimedia storage systems typically require large storage capacity, since the data rates for motion video are quite high. For example, 1.5 Mbps for MPEG-compressed video [6] corresponds to 1.5 GB for each 2-h movie.

A common application of multimedia servers is video-on-demand (VOD) [11, 13] service. In a VOD system, subscribers can choose both the video they wish to view and the time they wish to view it. This contrasts with services in which users can choose only from a small set of selections and/or watch them at prespecified times. For the purpose of this paper we consider VOD systems in which the service is homogeneous. In other words, a large number of multimedia streams of the same format, for example MPEG compressed, are stored and retrieved.

Pause and resume operations are among the most commonly used features on VCRs. In a VOD environment, there are inevitably "hot" videos which are requested by many viewers. Batching multiple viewing requests for the same video can greatly increase the number of viewers supportable by the video server. However, the requirement that each viewer be able to independently pause the video at any instant and later resume the viewing has caused difficulties in batching the viewers of each showing. In the VOD environment, viewers may be forced to wait before stream capacity becomes available to start the showing. However, once viewers have been admitted into the showing, it is generally felt undesirable to force them to incur a long wait for stream capacity to resume after a pause. To make sure that the pausing viewers can resume at any time, conventional approaches provide one video disk stream for each video request. Thus no batching is permitted. For each VOD server, there is a maximum number of video disk streams that can be supported. This is referred to as the *stream capacity* of the server. For example, if each disk can support $N_{vstream}$

simultaneous video streams, a video server with N_{disk} disks will have a stream capacity of $N_{vstream}N_{disk}$. (Note that $N_{vstream}$ will depend upon the disk arm scheduling algorithm and the amount of buffer available to support each round of retrieval [9, 15].) Clearly, the number of viewers supportable under the conventional approach is at most equal to the stream capacity of the server.

In this paper, we propose an efficient mechanism to support the pause and resume feature while allowing batching of concurrent viewers of the same video. The objective is to support a larger number of viewers than the stream capacity by employing the new concepts of *look-ahead stream scheduling* with *look-aside buffering*. This is referred to as the LASS scheme. Since VOD systems are often disk arm limited and the price of memory typically decreases far more rapidly than the price of disks, we explore techniques using memory buffering to reduce the disk arm requirements or to increase throughput (the maximum viewers supportable) for a given stream capacity. The intent is to allow batching of concurrent viewers and avoid backing up each viewer by a real video stream as much as possible. Instead, streams which are currently being used for other showings but are close to completion are used to back up batched viewers, so that pause and resume can be done at any time. These are referred to as look-ahead streams. When a stream designated as a look-ahead stream for another showing completes, if the backed up viewer is not in pause mode, we look for another look-ahead stream to replace it so the completing stream can be used instead to accommodate a showing for other waiting viewers. (Thus the identity of a viewer's look-ahead stream changes over time.) Otherwise, the completing stream will be used to support the resume of the pausing viewer.

We note that before the look-ahead stream becomes available, pausing and resuming have to be supported by the original stream through (look-aside) buffering of the missed content. That is, starting at the pausing point, the missed contents of the video are kept for the viewer in a look-aside buffer. Further details are given in Sect. 2. Real stream capacity will be reserved and put aside only if look-ahead streams are not available. An algorithm is provided to dynamically select the look-ahead streams and manage the buffer. Since all stream and buffer capacities reserved are put into common pools to be allocated on demand when the viewers go into pause mode, a substantial saving in hardware resources can be achieved. The scheme also provides a means to trade off system throughput and resume delay for a fixed amount of buffer. We develop a detailed simulation model to study the performance of the LASS scheme and conduct a sensitivity analysis of its buffer requirements. It is found that with sufficient buffer LASS can substantially improve the throughput of the video server compared to the conventional approach with no batching of viewers. Furthermore, *for a given amount of buffer, the improvement in throughput grows more than linearly with the stream capacity of the server*. That is to say, the LASS scheme operates with good economy of scale because it is easier to form look-ahead streams for video servers with larger stream capacity.

We briefly comment on some related work to support other aspects of VOD. Significant results were presented in [9, 13] regarding admission control techniques and the choice of service size to support multimedia applications. In [10], the concept of wait tolerance is explored to improve the batching effectiveness on viewer scheduling which determines the order that the viewer requests are scheduled. The proposed LASS scheme can work with any admission policy and viewer scheduling scheme. It provides an efficient way of utilizing disk stream capacity to support the scheduling of batched viewers with buffer. Certainly, LASS will be more effective if the admission policy and viewer scheduling scheme result in more and larger batches of viewers. The issues associated with supporting fast-forward and rewind are addressed in [2, 4]. Various papers have studied disk scheduling issues. In [15], a new formulation for disk arm scheduling schemes called grouped sweeping scheduling is proposed and analyzed. The goal is to minimize the buffer requirement. A similar concept is considered in [5]. Furthermore, [12] studies storage management and disk access algorithms in a disk array environment using this grouping approach. In [1, 13, 3], the issue of scalable data placement is considered. A somewhat different scheduling and placement perspective is taken in [14]. Nevertheless, both disk scheduling and video placement, each of which can affect the stream capacity of the VOD server, are orthogonal to the purpose of this paper: we assume a given stream capacity for the VOD server and try to maximize the number of viewers supportable by the server.

In Sect. 2, the proposed LASS scheme is presented. Section 3 contains a mathematical model and analysis of LASS. We describe the simulation model, assumptions and performance results in Sect. 4. Concluding remarks are provided in Sect. 5.

2 The LASS scheme

The preliminary concept of the LASS scheme is explained in Sect. 2.1. We also describe the stream and buffer states and provide examples. In Sect. 2.2 we present the stream scheduling algorithm where it is guaranteed that a paused viewer will be able to resume without incurring any delay. In Sect. 2.3 we describe an extension of the algorithm to provide a means to control the probability of delayed resume without strictly guaranteeing the capability to immediately resume. This offers a trade-off between the quality and the cost of the VOD service. The pause and resume operations are discussed in Sect. 2.4. Finally, we consider stream and viewing completions in Sect. 2.5.

2.1 Preliminaries

2.1.1 Basic concept

As mentioned above, the intent of look-ahead scheduling is not to back up each viewer with a real stream capacity so he can pause and resume at any time, but to back viewers up with a stream which is currently being used for another showing to be completed in the near future. This stream will be referred to as a *look-ahead* stream. We note that a look-ahead stream cannot simultaneously back up multiple viewers. Thus a stream can only be the look-ahead stream

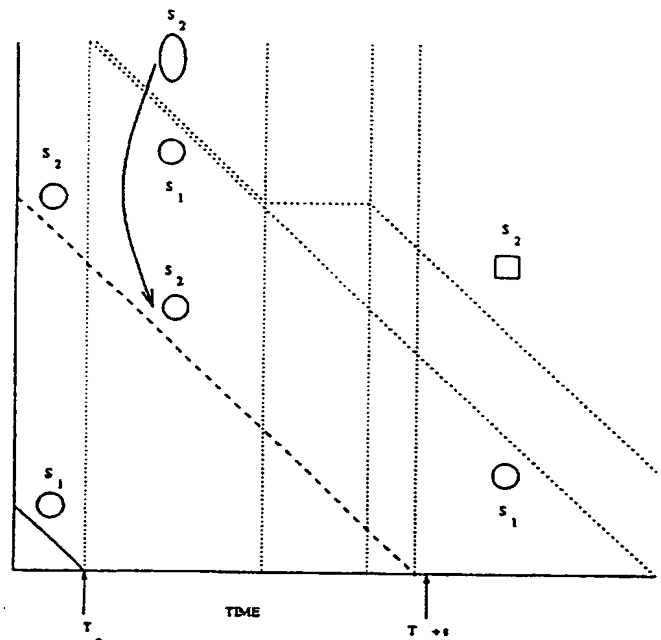

1

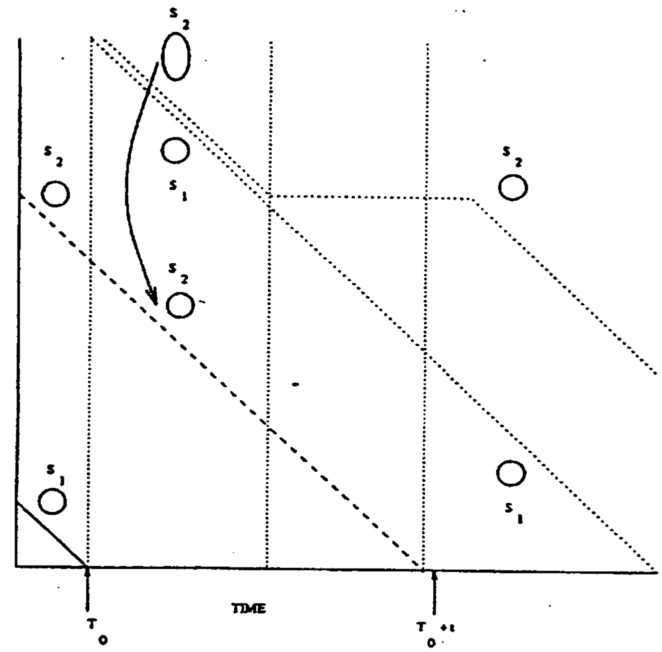

2

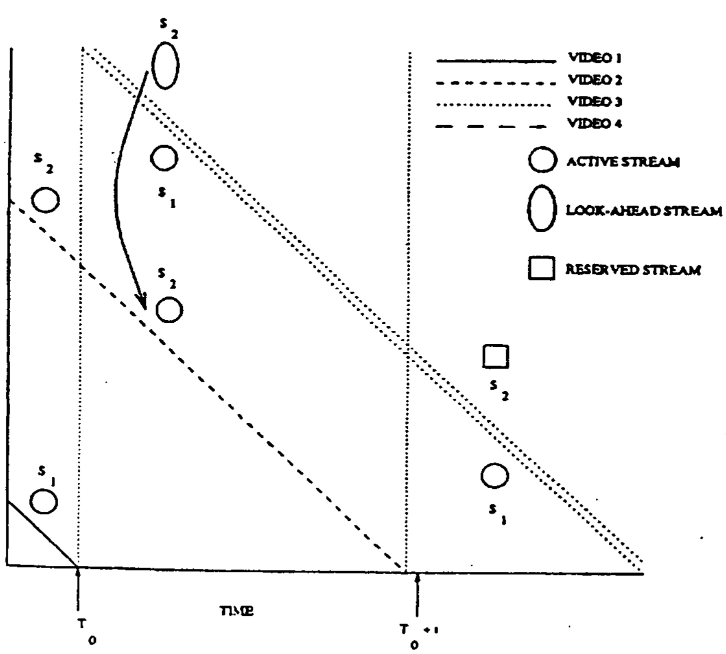

3

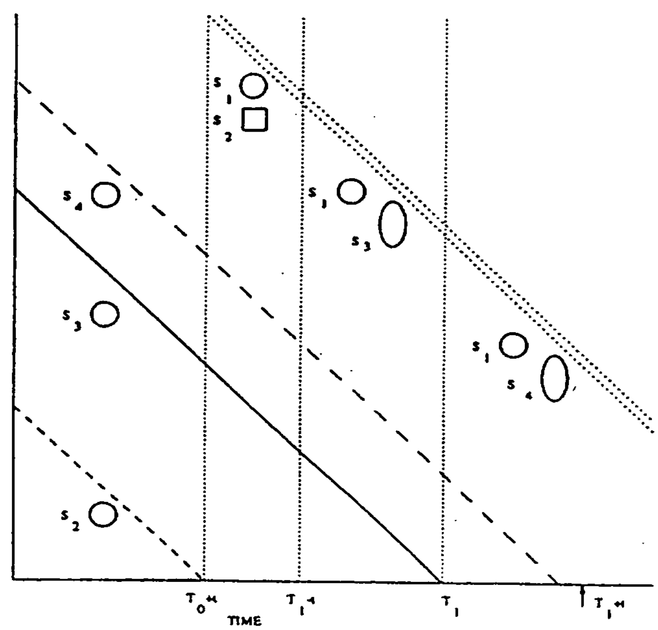

4

Fig. 1. Short pause scenario
Fig. 2. Long pause scenario
Fig. 3. Scenario with no pauses
Fig. 4. Scenario with additional future completions. Symbols correspond to Figs. 1–4

for one viewer at a time. When multiple viewers are batched together, only a single stream will actually be used to show the video. This stream will be referred to as the *active* stream.

Consider a scenario in which a VOD server has a buffer capable of storing t time units of a currently showing movie. We can let two viewers share the same video stream as long as another stream will become available within an interval of t time units. This is referred to as the *look-ahead interval*. Before the look-ahead stream becomes available, pausing and resuming have to be supported by the original stream through (look-aside) buffering of the missed content. This approach eliminates the need for a real stream capacity during the look-ahead interval. Look-aside buffering can also save subsequent stream requirements for short pauses. For

a pause which is less than the look-ahead interval, the resuming viewer does not need to use a new stream. Though lagging behind the showing of his original stream, the viewer can get the video contents from the buffer.

Thus the look-aside buffering saves stream capacity in two ways. First of all, it avoids putting aside stream capacity to support potential future pause/resume, if look-ahead streams can be found. Secondly, for short pauses the need for a new stream for the resuming viewer is eliminated. If there is not enough buffer to support look-ahead scheduling, a stream capacity is reserved to support future pause/resume. This stream will be referred to as a *reserved* stream. When a reserved stream is allocated, the number of useable streams of the VOD server is reduced by one. With a reserved stream, the (concurrent) viewers can pause at any time; when resuming, the reserved stream becomes the active stream to be viewed.

Example 1. Consider the scenario illustrated in Fig. 1. Each line shows the progression of a particular video request. (The horizontal axis indicates the time and the vertical axis indicates the amount of video remaining to be shown. Thus a horizontal line segment indicates a pause.) We indicate the active, look-ahead and reserved streams required to support each batched showing. At time T_0, the stream S_1 showing video 1 completes. Assume that the stream S_2 showing video 2 is to be completed at a time less than or equal to $T_0 + t$. Two requests for video 3 are at the head of the wait queue. Since sufficient buffer is available to support t time units, the two requests can be batched to share a common active video stream S_1. Stream S_2 will be used as the look-ahead stream. Suppose now that one of the new viewers pauses before $T_0 + t$. If it is a short pause (specifically, a pause of less than t units of time), the viewer can resume and be supported by the look-ahead buffer. After stream S_2 completes, it temporarily becomes a reserved stream. We will discuss later the way in which S_2 can be released.

The corresponding long pause scenario (longer than t units of time) is shown in Fig. 2. Here the resuming viewer can be supported by S_2, which becomes available before he resumes.

Note that at the time the video shown by a look-ahead stream completes, if another playing or reserved stream can be found which will be completed within the look-ahead interval, a new look-ahead stream can be designated and the completing stream can be used to schedule other viewers. So a viewer may be supported by a sequence of *different* look-ahead streams over time.

We now go back to Example 1: If neither of the two viewers of video 3 pauses before $T_0 + t$, S_2 will become a reserved stream to support S_1 (as shown in Fig. 3) until another stream can be found to serve as the look-ahead stream. In Fig. 4, stream S_3 of another showing of video 1 will be completed at T_1 and stream S_4 of video 4 will be completed before $T_1 + t$, assuming no pauses. Then, with sufficient buffer, S_3 can be used as the new look-ahead stream for S_1 after $T_1 - t$ and stream S_2 can be released to schedule other waiting requests. At time T_1, the look-ahead stream can be switched to S_4 and the previously designated look-ahead stream S_3 can be released to schedule waiting viewers.

Table 1. Data structure to maintain stream status

Stream ID	Active	Reserved	Look-ahead	Video ID
1	y		4	1
2		y	5	
3		y	5	
4	y			2
5	y			3
6		y		
7				
8				
9				
10				

Each member of a group of concurrent viewers of the same stream has to be supported by either the real stream showing the video or some backup stream, which can either be a look-ahead stream or a reserved stream. Each real or reserved stream can become a look-ahead stream of another showing. There is an additional level of complexity due to the fact that the viewer of a look-ahead stream may pause, so that the actual finishing time can be uncertain. To get around this problem, a stream chosen as a look-ahead stream cannot be allowed to pause. If the viewer pauses, the stream retrieval will continue and be buffered so that when the viewer resumes he will be able to view the video from the buffer. Once a viewer can obtain the remaining contents of the video from the buffer, no further stream capacity will be required for the video. Note that the buffer contents will not be released until the viewing is completed.

2.1.2 Remarks

First consider the state of the (look-aside) buffer. Each buffer block can be in one of three states – reserved, in-use, or available. As will be explained in detail later, during the scheduling of videos, certain portions of the buffer may be put into a reserved state to support pause/resume. A reserved buffer block changes into in-use when some video stream is stored into it. The remainder of the buffer is neither reserved nor in use, and is therefore available for future allocation.

We now examine the data structure used to keep track of the stream status. The VOD server has an upper limit on the number of concurrent video streams supportable. A stream is considered to be *active* if it is supporting an actual showing of a video. It is in *reserved* state if it is reserved to support future pause/resume of the concurrent viewers of some showing. If a stream capacity is neither active nor reserved, it is available for future showing. In Table 1 we show one way of doing the bookkeeping. The status of each stream (active, reserved or available) is recorded. If a stream is designated also as a look-ahead stream for a viewer in another showing, information on that showing stream is provided in the *look-ahead* field. We note that at scheduling time, we only associate the look-ahead streams with the active showing stream. This is equivalent to associating them with the set of viewers watching that video. The mapping of streams to specific viewers is only done at pausing time, as explained in Sect. 2.4.

Example 2. Assume that three requests for video 1 are scheduled at T_0 and at that time there are no active streams to be

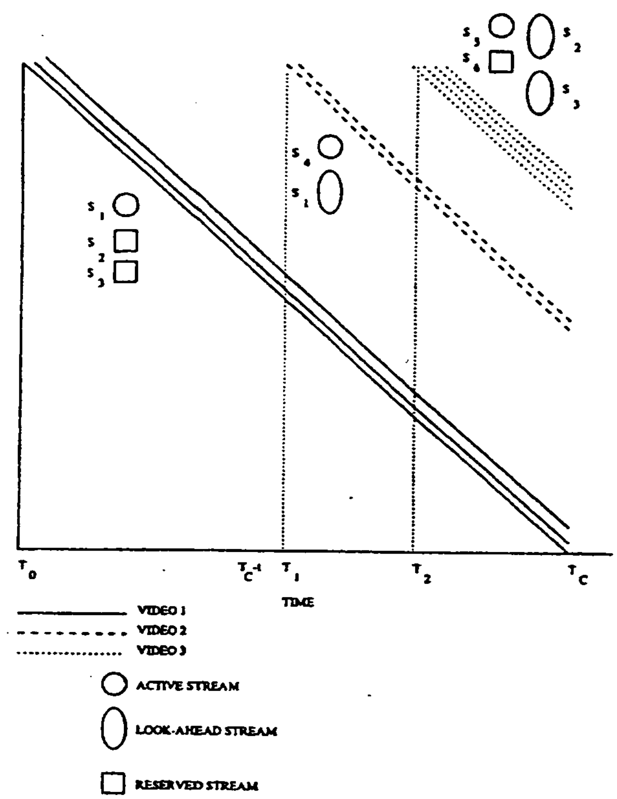

Fig. 5. Stream assignment under batching

completed in the next t units of time. Stream S_1 is chosen as the active stream, and streams S_2 and S_3 are designated as reserved streams for the concurrent viewers of S_1 (see Fig. 5 and Table 1, reserved field of S_2 and S_3.) We note that these reserved streams can be released if some other look-ahead stream becomes available later on. The completion time of S_1 is T_C. At time $T_1 \geq T_C - t$, two video requests for video 2 get scheduled. S_1 is now within t units of time to completion, and we assume there is sufficient buffer to support S_1 as a look-ahead stream. We can choose S_4 as the active stream and use S_1 as the look-ahead stream (see Table 1, *look-ahead* field of stream S_1). We note that the viewers of video 2 are not viewers of S_1, but merely use S_1 as a look-ahead stream to support future pause and resume operations. If an additional four requests for video 3 get scheduled at time T_2 between T_1 and T_C, the stream S_5 can be used as the active stream and the streams S_2 and S_3 as the look-ahead streams, assuming sufficient buffer. In addition, S_6 is needed as a reserved stream (see Table 1 in the look-ahead fields of S_2 and S_3 and the reserved field of S_6). Table 1 shows the stream status at this point, where there are nine viewers consuming six stream capacities.

2.2 *The scheduling algorithm*

Assume that the VOD server has a look-aside buffer of size B and a stream capacity of N_{MAX}. Let N_{RESRV} be the number of reserved streams in the system and N_{ACT} be the number of active streams showing videos. Let B_{RESRV} be the amount of look-aside buffer reserved and B_{USE} be the amount of look-aside buffer currently in use. We further assume that for each unit of time a showing requires K Mbytes of data.

If N_W customers are waiting for a particular video when the video is selected for showing, the following procedure determines the largest number C of viewers that can be scheduled to allow for pause/resume. Our approach uses as many look-ahead streams as possible, given the buffer constraint, and supports the remaining viewers with reserved streams.

More specifically:

1. We first need to determine the maximum number of additional look-ahead streams supportable given the current buffer usage. This is referred to as N_{LAHEAD} and is the minimum of the following two quantities:
   - The number of video streams (not yet marked as look-ahead streams) to be completed in the next t units of time assuming no pausing, where t is the length of the look-ahead interval, a prespecified operating parameter. These are the potential look-ahead streams.
   - The number of additional look-ahead streams supportable by the current state of the buffer.

   Let us order the potential look-ahead streams based on their remaining time to completion. It is obvious that from a buffering viewpoint we should choose look-ahead streams based on this order. Assuming that the ith potential look-ahead stream has the remaining time $\alpha_i t$ until completion, it will need a buffer of size $K\alpha_i t$ to be reserved if chosen. This buffer amount is needed for saving the video contents (until completion) if the current viewer of this potential look-ahead stream transfers into a pause state. (A buffer of size $K\alpha_i t$ guarantees that the rest of the showing can be streamed into the buffer, even in the case of immediate pause). If x look-ahead streams are chosen, an amount of $xK\alpha t$ additional reserved buffer will be needed to handle pausing of their associated viewers, where αt is the average remaining time to complete for the first x potential look-ahead streams, i.e $\alpha = \frac{\sum_{i=1}^{x} \alpha_i}{x}$. In addition, an amount of Kt buffer needs to be reserved to support short pausing of the new group of viewers (currently waiting to be scheduled) before the look-ahead streams become available. Hence, with x look-ahead streams chosen, the total amount of buffer that needs to be reserved is $(Kt + xK\alpha t)$. Thus, from the buffer viewpoint the maximal number of supportable look-ahead streams is the largest value x such that the buffer constraint is satisfied. Denote this maximum number of supportable look-ahead streams by N_B.
2. If $N_{LAHEAD} \geq N_W - 1$, we can schedule all these requesting viewers with one real stream for showing the video and $N_W - 1$ look-ahead streams. In this case C will equal N_W. Otherwise, the number of look-ahead streams used is N_{LAHEAD}. In order to schedule additional viewers not backed up by the look-ahead streams, we would need to obtain some stream capacities to be put into reserved state. The reserved streams obtainable must be smaller than the number of streams avail-

able, N_{AVAIL}, which is equal to $N_{MAX} - (N_{RESRV} + N_{ACT})$. If $N_{AVAIL} \geq (N_W - N_{LAHEAD} - 1)$, so that $N_W - N_{LAHEAD} - 1$ or more streams are available, then all requested viewers can be scheduled using $N_W - N_{LAHEAD} - 1$ reserved streams. Again, we have C equal to N_W. Otherwise, C will be $N_{AVAIL} + N_{LAHEAD}$.

Let D be the number of look-ahead streams used. We will then set B_{RESRV} to $Kt+DK\alpha t+B_{RESRV}$ and increase N_{ACT} by one. Also if reserved streams are used, we need to increase N_{RESRV} accordingly. Furthermore, B_{USE} will be incremented when the reserved buffers are actually in use to support a pause action. (B_{RESRV} will be decremented by the same amount.) This buffer will be released when not needed.

We note that by reserving additional amounts of buffer, the look-ahead streams assigned can be stretched out further. For example, for an additional Kt amount of buffer that can be reserved within the next t units of time, we can allow the look-ahead streams to be available t time units later. This rule can be applied repeatedly.

When a stream designated as a look-ahead stream is completed, we check whether another stream can replace it as a look-ahead stream (i.e., within the look-ahead interval of t time units to completion), in which case the previous look-ahead stream becomes available and new viewer requests can be scheduled using its capacity. Otherwise the completing stream becomes a reserved stream. If another stream will be completed within $t + w$ units of time, the reserved stream can be replaced by that look-ahead stream after w units of time. It can then be used to schedule another showing. Further optimization can improve the throughput by allowing a resuming viewer to merge with a later-showing real stream, assuming the timing is compatible. Still, an appropriate look-ahead stream is required as before to support additional pausing in the future.

2.3 Extensions of the scheduling algorithm

The scheduling procedure described above guarantees that a pausing viewer will always be able to resume at any time without incurring delay. This is certainly a very desirable level of service from the perspective of the viewer. However, from the perspective of the server it may be more cost-effective to tolerate a small probability of delaying viewers' resume requests (even if it means not charging these viewers). This could lead to a significant increase in the number of concurrent viewers supportable for a given hardware configuration or, conversely, a decrease in the hardware configuration required for a given throughput.

We note that in the above scheduling algorithm the buffer and the stream capacities reserved are treated as a pool to be shared by all viewers. They are not pre-allocated to specific viewers, and are allocated only on demand to pausing viewers. In other words only the viewers who actually pause will consume these buffer or stream capacities. Since in reality not all viewers are likely to pause at the same time, we run only a small risk of delaying the resuming viewers by reserving smaller amounts of streams and buffers.

Recall that the original buffer constraint on supportable look-ahead streams, x, is to find the maximum x such that $(Kt + xK\alpha t)$ is less than the amount of available buffer $(B + B_{RESRV} - B_{USE})$. We revise the buffer constraint above to be $\theta(Kt + xK\alpha t + B_{RESRV}) < (B - B_{USE})$, where θ, referred to as the buffer reservation ratio, is a tuning parameter. To avoid delay on resume, we need to set $\theta = 1$. Since not all viewers are pausing simultaneously, θ can be set to a lower value while maintaining a low probability that a resumed viewer will need to wait. Similarly, N_{AVAIL} can be redefined to be $N_{MAX} - (\theta' N_{RESERV} + N_{ACT})$, where θ', referred to as the stream reservation ratio, is a second tuning parameter. In the case of guaranteed no-delay resume, both θ and θ' are set to 1.

Note that B_{RESRV} and N_{RESRV} represent the amount of system resource being put aside to support fast resume operation. Lowering the value of θ (respectively, θ') reduces B_{RESRV} (respectively, N_{RESRV}). By reserving a smaller amount of buffers or stream capacities, the throughput can be increased at the expense of resume delay. In the performance result section we will show the sensitivities of θ and θ' values to the server throughput and the average resume delay or the probability of a no-delay resume. Certainly, a given level of performance in terms of throughput and resume delay/probability may be achieved by different combinations of θ and θ' values. This can have an implication on the video server configuration selection of whether to put in more memory buffers (so a larger amount of buffers can be put aside to support pause/resume) versus disk arms (so a larger amount of stream capacities can be put aside to support pause/resume). Thus the issue of optimal configuration selection and operating parameter (θ and θ') setting to achieve a given level of performance requires an analysis of the cost/performance trade-offs between buffer and disk arm. This is beyond the scope of this paper.

2.4 Pause and resume operations

We now examine how pause and resume operations are supported. For a set of C concurrent viewers sharing a common showing stream, there are an additional $C - 1$ backup streams, look-ahead or reserved, associated with the stream. The mapping of streams to specific viewers is only done at pausing time. To support pausing viewers we first consider using look-ahead streams, then reserved streams, and finally the active showing stream. For the next viewer issuing a pause request, we check whether there are enough look-ahead streams among the $C-1$ backup streams to cover all current pausing viewers and the newly requesting one. If so, we say the requesting viewer is supportable by a look-ahead stream. The active showing stream is affected only when C equals 1, or when all other viewers are already in pause mode.

If the viewer can be covered by a look-ahead stream, the reserved buffer is put into use to temporarily buffer the missed contents for the pausing viewer up to the length of the look-ahead interval t. When the pausing period exceeds t, the buffer is released if no other viewers are using it.

If the viewer cannot be supported by a look-ahead stream, he must be supported by either the active stream showing the video or a reserved stream. It is further checked whether the supporting stream is marked as a look-ahead

stream. If this is the case, the video stream will continue streaming the video contents into the buffer until completion. The operations associated with stream completion will be discussed in Sect. 2.5. In case the viewer is supported by an actual showing stream not marked as a look-ahead stream, the stream can be stopped and resumed later.

Next consider the resume operation: Upon receiving the resume request, a check is made as to whether the resuming point is available in the buffer. If so, the viewer resumes viewing from the buffer. Otherwise, a reserved stream will be converted into an actual showing stream to support the viewer.

2.5 Stream and viewing completions

We distinguish between stream completion and viewing completion. *Viewing* completion means that the viewer has completed the viewing of the video. *Stream* completion means the video stream has run to its end while the viewer may be in pause mode and not yet finished the viewing. In that case, the content of the video stream is saved in a buffer for the viewer. Hence, stream completion can occur either before or simultaneously with the viewing completion.

Let us examine the details of the stream completion operation. When a video stream is completed, it is checked whether this stream or any other associated completing reserved stream has been marked as a look-ahead stream. For each stream marked as a look-ahead stream, we need to find out whether another stream can be identified to replace it, and thus be switched into a look-ahead stream. This is addressed in detail later. If another stream can be switched into a look-ahead stream, that stream is designated as the new look-ahead stream to replace the completing video stream, and the completing stream is released as an available stream. The process of scheduling new video requests can then be initiated for the waiting video requests. Otherwise, if no other stream can be switched into a look-ahead stream, the completing stream becomes a reserved stream.

Next consider the viewing completion operation. As mentioned above, viewing completion may occur later than stream completion, since during a pause state for a viewer the video stream may continue and be saved in the buffer. Upon viewing completion, all buffers in use or reserved for the completing viewer are released if not needed by other viewers. We then check whether stream completion occurs at the same time. If so, the appropriate actions for stream completion described above are performed.

Finally, we discuss the process of dynamically switching look-ahead streams. Let ϵ be the lag (due to prior pausing) of the viewer associated with the look-ahead stream to the actual showing stream. If the value of ϵ is equal to zero, the look-ahead interval is set to t. Otherwise, the amount of available buffer is examined. Consider the case where there is sufficient buffer for an additional allocation of $\theta k \epsilon$. We will reserve that additional amount of buffer, and the look-ahead interval is set to t. For the case where there is insufficient buffer for an additional allocation of $\theta k \epsilon$, no additional buffer is reserved and the look-ahead interval is set to $t - \epsilon$. It is then checked whether any stream not yet marked as a look-ahead stream can terminate in the look-ahead interval, assuming pausing does not occur. If so, the earliest terminating stream is chosen as the look-ahead stream.

3 Mathematical model and analysis

A dominant factor in analyzing the performance of the LASS scheme is the relative frequency with which viewers choose to pause and resume in the course of a single showing. We use a state representation for each of these modes as shown in Fig. 6: A scheduled viewer initially enters the Resume state (R), where he stays for a period of time of length T_R. Then he moves to one of two Pause states; with probability p he enters the Short_Pause state (P_S), where the length of stay is T_{P_S}, and with probability $(1-p)$ he enters the Long_Pause state (P_L), where the length of stay T_{P_L}. The random variables T_R, T_{P_S} and T_{P_L} are assumed to be exponentially distributed with means $1/\mu_R$, $1/\mu_{P_S}$ and $1/\mu_{P_L}$ respectively. The two pause states reflect two types of interruptions which may occur during a video showing: momentary interruptions are represented by the Short_Pause state, while very long interruptions cause a transition to the Long_Pause state. Typically, the Short Pause Fraction, given by p, is higher than $1/2$.

In the analysis below we further assume that the system resources are fully utilized due to a high arrival rate. This implies a small probability of finding an idle stream at any time.

Under the LASS scheme, a viewer can incur delay before he returns to resume state, when $\theta < 1$ or $\theta' < 1$. However, the dedicated scheme, which assigns a separate stream to each scheduled viewer, incurs no such delay. Thus in comparing the two schemes we define a time interval that includes only the total time spent by a viewer in either a resume or a pause state.

We now compute the expected throughput ratio provided by the LASS scheme as compared to the dedicated stream approach. Denote by L' the extended length of a video show (due to pause/resume operations by a single viewer). Then L' can be represented as the sum of a series of disjoint pause/resume segments, where each segment consists of the viewing and subsequent pausing periods. The exception here is the last segment which does not contain a pausing period. Assume that the viewing segments, V_i, $i = 1, \ldots, n$, have length $T(V_i)$. In other words, $L' = \sum_{i=1}^{n} T(V_i)$.

Assuming the length of a video is L, the expected value of L' can be computed as

$$E(L') = L(1 + (\mu_R - 1/L)[pE(T_{P_S}) + (1-p)E(T_{P_L})]$$

For the analysis of the LASS scheme, assume first that $\theta = \theta' = 0$, so that neither look-ahead streams nor buffers are reserved for scheduled viewers. A stream capacity of N_{MAX} means that N_{MAX} concurrent viewers can be supported in the dedicated scheme. This is the number of viewers that can be started or completed in an interval of length $E(L')$. The LASS scheme, on the other hand, can support with its buffer space an additional number of viewers. For a given look-ahead interval t, a resuming viewer may require a buffer space of up to Kt Mbytes to tag along a video stream. Thus for a total buffer size B, let $N_B(t) = \frac{B}{Kt}$ denote the number of additional concurrent viewers supportable if each of them

requests Kt Mbytes. Note that each viewer who completes his viewing with a buffer to tag along a video stream starts using the buffer only after his first pause. Since the system can support $N_B(t)$ such viewers concurrently and each of them consumes the buffer allocation only for $\frac{E(L')-1/\mu_R}{E(L')}$ fraction of the showing, the additional number of viewers supportable during an interval of length $E(L')$ is $\frac{N_B(t)E(L')}{E(L')-1/\mu_R}$. Hence, the expected throughput ratio for $\theta = \theta' = 0$ is given by

$$R_T(0,0) = \frac{N_{MAX} + \frac{N_B(t)E(L')}{E(L')-1/\mu_R}}{N_{MAX}} \tag{1}$$

We now generalize to the case where $0 \le \theta \le 1$ but $\theta' = 0$. In other words, each additional batched viewer on a video showing is allocated a reserved buffer with probability θ. Thus, viewers who use buffers to complete their showings either hold these buffers for the entire duration of the showing (an average of $E(L')$ time units) with probability θ, or starting after their first pause (an average of $E(L') - 1/\mu_R$ time units) with probability $1-\theta$. Hence, we have

$$R_T(\theta,0) = \frac{N_{MAX} + \frac{N_B(t)E(L')}{(E(L')-1/\mu_R)(1-\theta)+E(L')\theta}}{N_{MAX}} \tag{2}$$

For the most general case, where $0 \le \theta \le 1$ and $0 \le \theta' \le 1$, we observe that the number of additional viewers that can be supported by the LASS scheme is the minimum of (1) the additional number of viewers the system can support when $0 \le \theta \le 1$, $\theta' = 0$ [as given in (2)], and (2) the additional number of viewers that can be supported when $\theta = 0$ and $0 \le \theta' \le 1$. Since the first case has already been analyzed, we now consider the second case. Denote by N_t the number of stream completions within a time interval of length t. Given the assumption that the load is sufficiently high, a completing stream is used again immediately for another showing. To simplify our analysis we make the assumption that new viewers are scheduled deterministically at periodic intervals, with an average interval length of $E(L')/N_{MAX}$ time units. For a given look-ahead interval t, we have $N_t = tN_{MAX}/E(L')$. Compared to the dedicated scheme, the additional number of viewers schedulable under the LASS scheme is limited by N_t/θ', since each of these viewers requires a look-ahead stream allocated with probability θ'. Note that any viewer that gets look-ahead stream support but completes the viewing using a buffer (to tag along a video stream) consumes one look-ahead stream capacity on the average of $1/\mu_R$ time units, until his first pause. (As explained in Sect. 2, within this $1/\mu_R$ time period the identity of the look-ahead stream changes over time, i.e. the allocation of look-ahead stream capacity to a viewer results in a sequence of $1/t\mu_R$ look-ahead stream assignments, each lasting t units of time.) Hence, within a time interval of length L' there can be on the average $L'\mu_R$ iterations of look-ahead stream capacity allocations for additional viewer support, and

$$R_T(\theta,\theta') = \frac{N_{MAX} + \min\left(\frac{E(N_t)}{\theta'} \cdot E(L')\mu_R,\ \frac{N_B(t)E(L')}{(E(L')-1/\mu_R)(1-\theta)+E(L')\theta}\right)}{N_{MAX}} \tag{3}$$

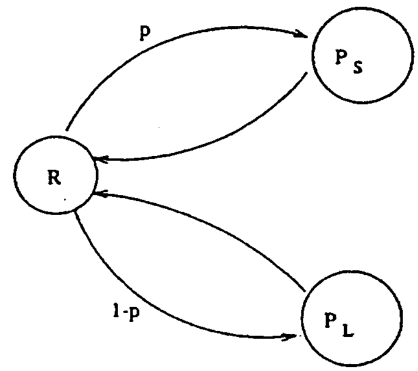

Fig. 6. Transition diagram for Pause/Resume simulation model

The above computation may be refined to incorporate the effect of allowing bounded viewer wait times on system throughput and the phenomenon that a fraction of the viewers who complete their viewings with dedicated streams start occupying these streams only after their first pause or later. Comparison of the analysis with simulation results is given in the next section.

4 Experimental results

4.1 Simulation model and assumptions

4.1.1 Video selection

Using the above pause/resume model, we simulated the arrivals of video requests by a Poisson process, with the interarrival time T distributed exponentially with mean $1/\lambda$.

Assume that there are M different videos available in the server. Upon arrival in the system, a viewer chooses to watch the ith video with probability p_i, $1 \le i \le M$. We assume that the request probabilities are homogeneous in time throughout the simulation. Obviously, any change in the distribution of hot videos can be reflected by choosing a simulation interval of different length.

The LASS performance was tested for a *Zipf* distribution [7]: $p_i = \frac{c}{i}$, where $c = \frac{1}{\sum_{i=1}^{M} \frac{1}{i}}$ is a normalization constant.

4.1.2 System model

Scheduling of new viewers is done periodically, depending on the window size which is given as a parameter.

Each arriving viewer is handled simultaneously by two subsystems:

1. In the Look-ahead subsystem we implement the LASS scheme as described in Sect. 2.
2. The No_Batching subsystem simulates the naive scheme, which assigns to each viewer a dedicated stream. Thus the maximal number of viewers served by the system at any time cannot exceed N_{max}, the total number of streams. An arriving viewer is scheduled if there is an

available stream. A pause is implemented by delaying the stream completion time by the length of the viewer's pause. Hence, a viewing completion is identified with a stream completion event. The stream is then released into the pool of available streams. Thus the stream allocation policy guarantees that a pausing viewer can always resume without waiting.

4.2 *Performance results*

We first describe the set of base parameter values used in the simulations. Assuming a video showing with the rate 1.5 Mbits/s, the amount of buffer reserved for a 1-min pause is $K = 11.25$ Mbytes. We assume that the number of videos is $M = 100$, and the length of a video is $L = 120$ min. We further assume that a viewer's pause may last on the average either $1/\mu_{P_S} = 1$ min with a short pause probability of $p = 0.7$, or else $1/\mu_{P_L} = 30$ min with a probability $(1-p) = 0.3$. The viewer remains in resume state on the average $1/\mu_R = 40$ min, an average of three pauses per viewing. The video request rate is nine arrivals per minute. (Note that to measure the server throughput we need to set the video request rate high enough to drive the video server to its full capacity. To prevent overflowing the system, we reject video requests based on the maximum number of viewers allowed into the system and the maximal lengths of the queues.) The length of the look-ahead interval was selected so as to maximize the throughput using a standard bracketing and bisection method [8].

In Fig. 7 we show the system maximal throughput versus server stream capacity for the case of $\theta = \theta' = 1.0$ compared to the No_batching scheme under two different buffer sizes. The maximum throughput ratio is defined to be the ratio of the maximum throughput of LASS to the throughput of the No_batching scheme. With $\theta = \theta' = 1.0$, the resume operations do not incur any wait time. Thus an arriving viewer is allocated a look-ahead stream and a look-aside buffer, which together serve as a backup for any pause which may take place during the viewer's video showing. It is important to note that for a given buffer size, the maximum throughput ratio increases with the server stream capacity. Hence the actual improvement in throughput, namely the additional number of viewers supportable by the system, grows more than linearly with the stream capacity of the server. This is due to the fact that it is easier to form look-ahead streams for video servers with larger stream capacity. In Table 2 we give the numbers of viewers scheduled per 120-min video showing for the scenarios of the LASS scheme shown in Fig. 7. Assuming a 10 GB buffer size, the additional number of viewers supportable per video showing (compared to the No_batching case) is 12 for $N_{max} = 100$, and 177 for $N_{max} = 1000$, respectively. This clearly shows that the LASS scheme operates with good economy of scale. The more than linear improvement in throughput is mainly due to two reasons. First of all, for large servers there are more batching opportunities. With a higher request rate, it is more likely that multiple requests for the same hot video will arrive in close time proximity. Secondly, for large servers the mean time between video completions is smaller. Hence it is more likely that a look-ahead stream will be found during a given look-ahead interval. In fact, a shorter look-ahead interval should be used as explained in Fig. 9.

Table 2. Number of viewers scheduled for increasing number of streams where $\theta = \theta' = 1.0$

No. streams/buffer(GB)	No_batching	10.0	14.0
100	86	98	100
400	343	397	411
600	511	602	620
800	678	812	837
1000	838	1015	1046

For the case $\theta = \theta' = 1.0$, Fig. 8 shows the sensitivity of the maximal throughput ratio to the length of the look-ahead interval, where $t = 2, \ldots, 16$ min, with $N_{max} = 1000$ and $B = 10$ GB. We show in Fig. 9 the length of the optimal look-ahead interval which maximizes the throughput for different buffer sizes with $N_{max} = 1000$ and 100, respectively. For a given server stream capacity, larger buffer size allows for a larger look-ahead interval. However, we observe that for a given buffer allocation, increasing the server stream capacity requires a smaller look-ahead interval, as stream completions occur more frequently.

Next we show the effect of varying θ and θ', which trades off service quality against further throughput improvement. Some paused viewers may need to wait to resume if θ and θ' are not set to 1. Figures 10–12 present the LASS performance with stream reservation ratios (θ') of 1, 0.5 and 0, where the buffer reservation ratio (θ) varies from 0 to 1. A buffer size of 10 GB and a stream capacity of 1000 are assumed. Figure 10 shows that the throughput ratio can be improved by reducing the buffer reservation ratio (θ) in addition to the stream reservation ratio (θ'). In Fig. 11, the average wait time is calculated only for viewers who need to wait for the resume. From Figs. 10, 11 and 12 we note that with a small wait probability (for example less than 4.9%) and an average wait time under 1 min, we can improve the throughput ratio from 1.21 to 1.36 (more than 12%), by setting $\theta = \theta' = 0$ rather than $\theta = \theta' = 1$.

In Fig. 13 we compare the simulation results with the analysis for $\theta' = 0$ and 0.5 as given in Eqs. 1 and 3, respectively. We assume a stream capacity of $N_{MAX} = 1000$ and a buffer size of 10 GB. For each value of θ, the same value of t that yields the maximal throughput was used for deriving the simulation and analysis results. The match between simulation and analysis is especially close for $\theta' = 0$.

In Figs. 14–17 we study the system behavior when θ and θ' are set to 0. Figure 14 presents the maximal throughput ratio versus the length of the look-ahead interval with $N_{max} = 800$ and $B = 10$ GB. Even for this case of $\theta = \theta' = 0$, the throughput is still affected by the length of the look-ahead interval, as it determines the maximum pause or amount of buffer allocation allowed for a resumed viewer to avoid another stream allocation. Furthermore, there is a trade-off between the maximal number of viewers the system can support concurrently (determined by the length of the look-ahead interval) and the average length of time a viewer needs to wait to resume. This gives rise to the optimization, used in our experiments, of the value of t that maximizes the throughput, while constraining the average wait time (to resume) to be bounded by 1 min. In Figs. 15, 16 and 17 we plot the maximal system throughput versus buffer

A variation of the LASS scheme with early stream release is considered, where we allow the release of paused streams: We increase the system throughput by allowing a stream to pause only up to a predetermined amount of time. Whenever a stream is paused for that amount of time, it is released and becomes available for the use of other viewers, and the corresponding paused viewer will need to get another stream when he resumes. In Fig. 18 we plot the improvement in throughput ratio of the LASS scheme with early stream release compared to the original scheme. The threshold to release the streams on pause is set to 1 min. Figure 19 shows the corresponding increase in wait probability. We note that an increase of almost 9% in throughput ratio incurs an increase of at most 2% in wait probability. As before, we allow the system throughput to increase as long as the average wait time incurred by the viewers waiting to resume does not exceed 1 min.

5 Summary

In a VOD environment there are invariably hot videos which are requested by many viewers. Batching multiple viewing requests for the same video can greatly increase the number of viewers supportable by the video server. However, the requirement that each viewer can independently pause the video at any instant and later resume the viewing without delay can cause difficulties in batching viewers for each showing. In this paper we propose an efficient mechanism to support the pause/resume feature while allowing batching of concurrent viewers on the same video. The proposed LASS scheme exploits the concept of *look-ahead stream scheduling* with *look-aside buffering*. It uses buffering to improve the number of concurrent viewers supportable. The concept of look-ahead scheduling is not to back up each viewer with a real stream capacity so he can pause and resume at any time, but rather to back it up with a (look-ahead) stream that is currently being used for another showing that is close to completion. Before the look-ahead stream becomes available, the pause and resume features have to be supported by the original stream through (look-aside) buffering of the missed content.

We have developed analytic and simulation models to study the performance of the proposed LASS scheme. We compare LASS with the conventional approach where no batching is allowed. It is found that with sufficient buffer LASS can provide a substantial improvement in throughput compared to the conventional approach. Furthermore, for a given amount of buffer, the improvement in throughput grows more than linearly with the stream capacity of the server. The LASS scheme operates with good economy of scale because it is easier to form look-ahead streams for video servers with larger stream capacity. LASS can also provide a means to trade off throughput and resume delay based on the amount of look-aside buffer. This trade-off is also analyzed.

References

1. Chang E, Zakhor A (1994) Scalable Video Data Placement on Parallel Disk Arrays. Proceedings IS&T/SPIE Symposium on Electronic Imaging – Conference on Image and Video Databases II. SPIE, February
2. Chen MS, Kandlur DD, Yu PS (1994) Support for fully interactive playout in a disk-array-based video server. Proceedings ACM Multimedia 94, San Francisco, Calif, October, pp 391–398
3. Chen MS, Kandlur DD, Yu PS (1995) Using rate staggering to store scalable video data in a disk-array-based video server. Proceedings Multimedia Computing and Networking 1995 (IS&T/SPIE Symposium on Electronic Imaging: Science and Technology). San Jose, Calif, February. SPIE 2417:338–345
4. Dey-Sircar JK, Salehi JD, Kurose JF, Towsley D (1994) Providing VCR capabilities in large-scale video servers. Proceedings ACM Multimedia 94, San Francisco, Calif, October, pp 25–32
5. Gemmell DJ (1993) Multimedia network file servers: multi-channel delay sensititive data retrieval. Proceedings ACM Multimedia 93, Anaheim, Calif, August, pp 243–250
6. International Organization for Standardization (1992) Coding of motion pictures and associated audios for digital storage media at up to 1.5 Mbits/s. IS 11172, November
7. Knuth D (1973) The art of computer programming, volume 3: Sorting and searching. Addison-Wesley, Reading, Mass
8. Press W, Flannery B, Teukolsky S, Vetterling W (1986) Numerical recipes. Cambridge University Press, Cambridge, UK
9. Rangan PV, Vin HM (1991) Designing file systems for digital video and audio. Proceedings, 12th ACM Symposium on Operating Systems
10. Shachnai H, Yu PS (1995) The role of wait tolerance in effective batching: a paradigm for multimedia scheduling schemes. IBM Research Report RC 20038, April
11. Sincoskie W (1991) System architecture for large scale video on demand. Comput Networks ISDN Syst 22:155–162
12. Tobagi FA, Pang J, Baird R, Gang M (1993) Streaming RAID – a disk array management system for video files. Proceedings ACM Multimedia 93, Anaheim, Calif, August, pp 393–400
13. Vin HM, Rangan PV (1992) Designing a multi-user HDTV storage server. UCSD Tech Rep CS92-225, January
14. Wolf JL, Yu PS, Shachnai H (1995) DASD dancing: a disk load balancing optimization scheme for video-on-demand computer systems. Proceedings ACM Sigmetrics 95, Ottawa, May
15. Yu PS, Chen MS, Kandlur DD (1993) Grouped Sweeping Scheduling for DASD based Multimedia Storage Management. Multimedia Systems 1:99–109

PHILIP S. YU received the BS degree in EE from National Taiwan University, Taipei, Taiwan, Republic of China, in 1972, the MS and PhD degrees in EE from Stanford University in 1976 and 1978, respectively, and the MBA degree from New York University in 1982. Since 1978 he has been with the IBM Thomas J. Watson Research Center. Currently he is manager of the Architecture Analysis and Design group. His current research interests include transaction and query processing, database systems, data mining, parallel and distributed processing, multimedia systems, disk arrays, computer architecture, performance modelling, and workload analysis. He has published more than 180 papers in refereed journals and conferences, and over 130 research reports and 80 invention disclosures. He holds or has applied for 25 US patents. Dr. Yu is a Fellow of the IEEE. He is an editor of *IEEE Transactions on Knowledge and Data Engineering*. In addition to serving as program committee member on various conferences, he has served as the program chair of the 2nd International Workshopon Research Issues on Data Engineering: Transaction and Query Processing and the program co-chair of the 11th International Conference on Data Engineering. He has received several IBM and external honors, including Best Paper Award, IBM Outstanding Innovation Award, Outstanding Technical Achievement Award, Research Division Award, and 13 Invention Achievement Awards.

JOEL L. WOLF received his PhD from Brown University in 1973 and his ScB from the Massachusetts Institute of Technology in 1968, both in Mathematics. He is currently a Research Staff Member at the IBM Thomas J. Watson Research Center, with interests in mathematical optimization. In 1988, he won an IBM Outstanding Innovation Award (OIA) for his work on the Placement Optimization Program, a practical optimization technique to solve the disk file assignment problem. In 1994, he won another OIA for his work on parallel query processing. He has also been an Assistant Professor of Mathematics at Harvard University, as well as a Distinguished Member of Technical Staff and manager at Bell Laboratories. Dr. Wolf is a senior member of IEEE, INFORMS, and ACM and the author of numerous papers and patents.

HADAS SHACHNAI received BS and PhD degrees in Computer Science from the Israel Institute of Technology (Technion) in 1986 and 1991, respectively. She joined the Department of Computer Science at the Technion as Assistant Professor in 1991. Currently she is a visiting postdoc at the IBM Thomas J. Watson Research Center. Her research interests are performance evaluation and probabilistic modeling of computing and multimedia systems.